Aim:

Review:

Goals:

- Direct solve BSS mat.
  → HBSS mat.
- Butterfly
- PDE solvers | compact op.

  theory

(Recursive) Coordinate Bisection (RCB)

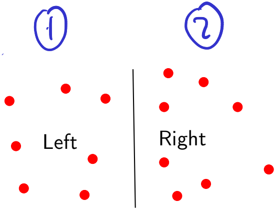

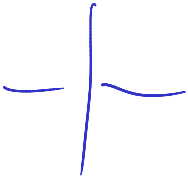

	1	2
1	full rank	low rank
2	low rank	full rank

Block-separable matrices

Particles:

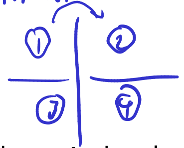

$$A = \begin{array}{cc} \overbrace{\qquad\qquad\qquad\qquad}^{n} & \\ \begin{bmatrix} D_1 & A_{12} & A_{13} & A_{14} \\ A_{21} & D_2 & A_{23} & A_{24} \\ A_{31} & A_{32} & D_3 & A_{34} \\ A_{41} & A_{42} & A_{43} & D_4 \end{bmatrix} & \left.\vphantom{\begin{bmatrix} 1\\1\\1\\1 \end{bmatrix}}\right\} n \end{array}$$

D_i: full rank

if sym: $A_{ij} = A_{ji}^T$

where A_{ij} has low rank: How to capture rank structure?

$$A_{ij} \approx \underset{O(n)}{P_i} \underbrace{(A_{ij})_{I,J}}_{\tilde{A}_{ij}} \underset{O(n)}{\Pi_j}$$

?

Optimistically: $k \times k$

k: off-diag interaction rank

Proxy Recap

Saw: If A comes from a kernel for which Green's formula holds, then the same skeleton will work for all of space, for a given set of sources/targets. What would the resulting matrix look like?

Rank and Proxies

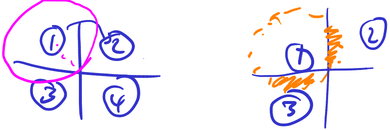

Unlike FMMs, partitions here do not include "buffer" zones of near elements. What are the consequences?

Idea:

- Add extra proxies from nearby particle groups
- Tolerate the increased rank.

Block-Separable Matrices

A *block-separable matrix* looks like this:

$$A \approx \begin{bmatrix} D_1 & P_1\tilde{A}_{12}\Pi_2 & P_1\tilde{A}_{13}\Pi_3 & P_1\tilde{A}_{14}\Pi_4 \\ P_2\tilde{A}_{21}\Pi_1 & D_2 & P_2\tilde{A}_{23}\Pi_3 & P_2\tilde{A}_{24}\Pi_4 \\ P_3\tilde{A}_{31}\Pi_1 & P_3\tilde{A}_{32}\Pi_2 & D_3 & P_3\tilde{A}_{34}\Pi_4 \\ P_4\tilde{A}_{41}\Pi_1 & P_4\tilde{A}_{42}\Pi_2 & P_4\tilde{A}_{43}\Pi_3 & D_4 \end{bmatrix}$$

$$C = \begin{pmatrix} 0 & \tilde{A}_{12} & \cdots \\ \tilde{A}_{21} & \ddots & \\ & & 0 \end{pmatrix}$$

For matvec: $\underbrace{D}_{N\times\sqrt{N}}\vec{x} + \underbrace{P}_{N\times K}\underbrace{C}_{N\times K}\underbrace{\Pi}_{K\times N}\vec{x}$

$$D = \begin{pmatrix} \boxed{D_1}^{\sqrt{N}} & \\ & \ddots \end{pmatrix}$$

Here:

- $\tilde{A}_{ij}$ smaller than A_{ij}
- D_i has full rank (not necessarily diagonal)
- P_i shared for entire row
- Π_i shared for entire column

Shape of Π_i:

$\Pi_i =$

$$\Pi = \begin{pmatrix} \Pi_1 & & \\ & \ddots & \\ & & \Pi_4 \end{pmatrix}$$

$$\tilde{x} = \Pi x$$

Block-Separable Matrix: Questions

Q: Why is it called that?

u

$A = \;$ | | | $\quad$ b_i

a_i

$u(x,y) = \sum_i a_i(x)\, b_i(y)$

"separation of variables"

Q: How expensive is a matvec?

$O(N^{3/2})$ by analogy to single-level Barnes-Hut.

assuming: $\sqrt{N}$ particles in a box/group

Q: How about a solve?

?

BSS Solve (I)

Separate out 'coarse' unknowns. Use the following notation:

$$B = \begin{bmatrix} 0 & P_1\tilde{A}_{12} & P_1\tilde{A}_{13} & P_1\tilde{A}_{14} \\ P_2\tilde{A}_{21} & 0 & P_2\tilde{A}_{23} & P_2\tilde{A}_{24} \\ P_3\tilde{A}_{31} & P_3\tilde{A}_{32} & 0 & P_3\tilde{A}_{34} \\ P_4\tilde{A}_{41} & P_4\tilde{A}_{42} & P_4\tilde{A}_{43} & 0 \end{bmatrix}$$

$A\vec{x} = \vec{b}$

and

$$D = \begin{bmatrix} D_1 & & & \\ & D_2 & & \\ & & D_3 & \\ & & & D_4 \end{bmatrix}, \quad \Pi = \begin{bmatrix} \Pi_1 & & & \\ & \Pi_2 & & \\ & & \Pi_3 & \\ & & & \Pi_4 \end{bmatrix}.$$

$A = D + B\Pi$

$$\begin{pmatrix} D & B \\ -\Pi & \mathrm{Id} \end{pmatrix} \begin{pmatrix} \vec{x} \\ \vec{\tilde{x}} \end{pmatrix} = \begin{pmatrix} \vec{b} \\ \vec{0} \end{pmatrix} \iff \begin{aligned} D\vec{x} + B\vec{\tilde{x}} &= \vec{b} \\ -\Pi\vec{x} + \vec{\tilde{x}} &= \vec{0} \end{aligned}$$

$\Leftrightarrow \Pi\vec{x} = \vec{\tilde{x}}$

BSS Solve (II)

Q: What are the matrix sizes? The vector lengths of $\boldsymbol{x}$ and $\tilde{\boldsymbol{x}}$?

$\Pi :$

Now work towards doing *just* a 'coarse' solve on $\tilde{\boldsymbol{x}}$, using the Schur complement:

$$\begin{pmatrix} D & B \\ -\Pi & \mathrm{Id} \end{pmatrix}\begin{pmatrix} \vec{x} \\ \tilde{x} \end{pmatrix} = \begin{pmatrix} \vec{b} \\ \vec{0} \end{pmatrix} \qquad I(\Pi D^{-1})\cdot \rightsquigarrow \left(\begin{array}{cc|c} \Pi D^{-1} D & \Pi D^{-1} B & \Pi D^{-1} \vec{b} \end{array}\right)$$

$$\oplus$$

$$\begin{pmatrix} 0 & \mathrm{Id} + \Pi D^{-1} B \end{pmatrix}\begin{pmatrix} \vec{x} \\ \tilde{x} \end{pmatrix} = \left(\Pi D^{-1} \vec{b} \right)$$

$$\left(\mathrm{Id} + \Pi D^{-1} B\right) \tilde{x} = \Pi D^{-1} \vec{b}$$

BSS Solve (III)

Focus in on the second row:

$$\circledast \qquad (\underbrace{\mathrm{Id}}_{\text{on-diag}} + \overbrace{\Pi D^{-1} B}^{\text{off-diag}})\widetilde{\boldsymbol{x}} = \Pi D^{-1}\boldsymbol{b}$$

Every non-zero (i.e. off-diagonal) entry in $\Pi D^{-1} B$ looks like

$$\Pi_i D_i^{-1} P_i \hat{A}_{ij}$$

Define a diagonal entry:

$$\tilde{A}_{ii} = \left(\Pi_i D_i^{-1} P_i\right)^{-1}$$

Idea: Hit $\circledast$ from left with block-diag$_i(\tilde{A}_{ii})$

BSS Solve (IV)

Next, left-multiply $(\mathrm{Id} + \Pi D^{-1} B)$ by block-diag$_i(\tilde{A}_{ii})$:

$$\begin{pmatrix} \hat{A}_{11} & & & \\ & \hat{A}_{22} & & \\ & & \hat{A}_{33} & \\ & & & \hat{A}_{44} \end{pmatrix} (\mathrm{Id} + \Pi D^{-1} B) = \underbrace{\begin{pmatrix} \tilde{A}_{11} & \tilde{A}_{12} & & \\ \hat{A}_{21} & \ddots & & \\ & & & \\ & & & \tilde{A}_{44} \end{pmatrix}}_{\tilde{A}}$$

$$(\mathrm{Id} + \Pi D^{-1} B)\widetilde{\boldsymbol{x}} = \Pi D^{-1} \boldsymbol{b}$$

$$\tilde{A}\tilde{x} = \text{block-diag}_i(\tilde{A}_{ii})\, \Pi D^{-1} b$$

$$\Pi \hat{x} = \tilde{x} \qquad \cancel{\hat{x} = \Pi^{-1}\tilde{x}} \qquad D\hat{x} + B\hat{x} = b$$

BSS Solve: Summary

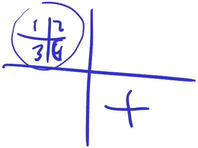

What have we achieved?

- Instead of solving a linear system of size
$$(N_{L0\,\text{boxes}} \cdot m) \times (N_{L0\,\text{boxes}} \cdot m)$$
we solve a linear system of size
$$(N_{L0\,\text{boxes}} \cdot K) \times (N_{L0\,\text{boxes}} \cdot K),$$
which is cheaper by a factor of $(K/m)^3$.
- We are now only solving on the skeletons.

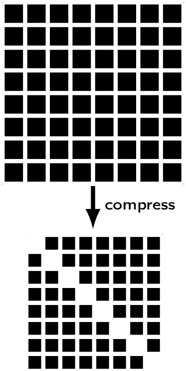

(Figure following G. Martinsson, drawn by A. Fikl)

Hierarchically Block-Separable

To get to $O(N)$, realize we can *recursively*

- group skeletons
- eliminate more variables.

Where does this process start?

Demo: Skeletonization using Proxies (Hierarchical)

Hierarchically Block-Separable

In order to get $O(N)$ complexity, could we apply this procedure recursively?

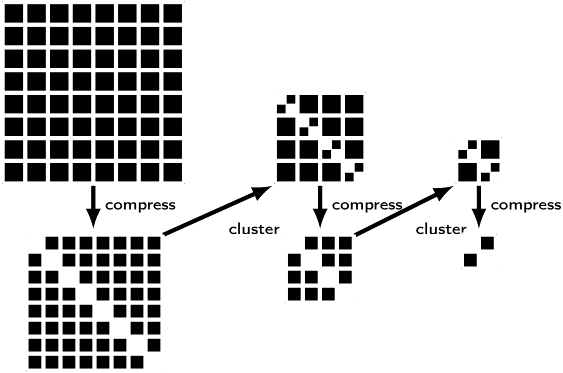

(Figure following G. Martinsson, drawn by A. Fikl)

Hierarchically Block-Separable

- Using this hierarchical grouping gives us *Hierarchically Block-Separable* (*HBS*) matrices.
- If you have heard the word $\mathcal{H}$*-matrix* and $\mathcal{H}^2$*-matrix*, the ideas are very similar. Differences:
    - $\mathcal{H}$-family matrices don't typically use the ID (instead often use *Adaptive Cross Approximation* or *ACA*)
    - $\mathcal{H}^2$ does target clustering (like FMM), $\mathcal{H}$ does not (like Barnes-Hut)

Outline

Recap: Fast Fourier Transform

The *Discrete Fourier Transform (DFT)* is given by:

$$X_k = \sum_{n=0}^{N-1} x_n e^{-\frac{2\pi i}{N} nk} \quad (k = 0, \ldots, N-1)$$

The foundation of the *Fast Fourier Transform (FFT)* is the factorization:

$$X_k = \underbrace{\sum_{m=0}^{N/2-1} x_{2m} e^{-\frac{2\pi i}{N/2} mk}}_{\text{DFT of even}-\text{indexed part of } x_n} + e^{-\frac{2\pi i}{N} k} \underbrace{\sum_{m=0}^{N/2-1} x_{2m+1} e^{-\frac{2\pi i}{N/2} mk}}_{\text{DFT of odd}-\text{indexed part of } x_n} .$$

FFT: Data Flow

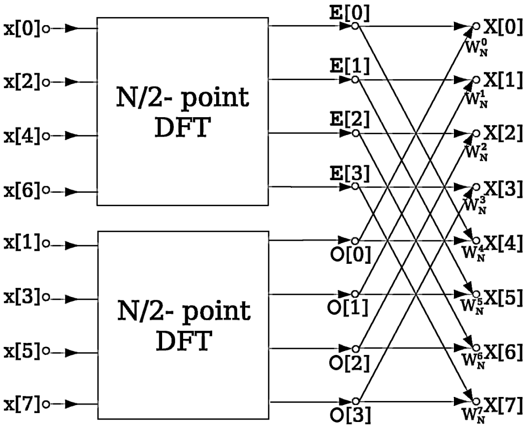

(Figure credit: Wikipedia)
Perhaps a little bit like a butterfly?

Fourier Transforms: A Different View

$$\mathcal{F}f(x) = \int e^{ixt} f(t)\, dt$$

Claim:

> *The [numerical] rank of the normalized Fourier transform with kernel $e^{i\gamma xt}$ is bounded by a constant times γ, at any fixed precision ϵ.*

(i.e. rank is proportional to the area of the rectangle swept out by x and t)
[O'Neil et al. '10]

Demo: Butterfly Factorization (Part I)

Recompression: Making use of Area-Bounded Rank

How do rectangular submatrices get expressed so as to reveal their constant rank?